WEST Search History

Hide Items | Restore | Clear | Cancel

DATE: Thursday, May 18, 2006

Hide?	Set Name	Query	Hit Count
		DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ	
☐	L27	L26 and (default near5 buffer$1)	1
☐	L26	L25 and subnet	93
☐	L25	L24 and infiniband	289
☐	L24	L23 and (channel near5 adapter$1)	289
☐	L23	buffers and infiniband	1031
☐	L22	L21 and infiniband	0
☐	L21	L20 and queue	4
☐	L20	L19 and buffer$1	4
☐	L19	L18 and cluster$1	4
☐	L18	('6978300' |'6938138' |'6799220' |'6747949')!.ABPN1,NRPN,PN,TBAN,WKU.	8
☐	L17	L16 and (buffer near5 usage)	8
☐	L16	(message$1 and buffer$1 and switched and fabric and queue and pair$1 and threshold and value) and @py<=2002	141
☐	L15	(general services agent) and buffer$1 and cluster$3 and switched and fabric and queue and message$1 and @py<=2002	0
☐	L14	(buffer$1 and cluster$3 and switch$2 and fabric and message$1 and ((queue near5 pair$1) or qp$1) and threshold and value$1) and @py<=2002	1
☐	L13	(host and buffer$1 and cluster$1 and queue and threshold and post and pre and incom$3 and data and message$1 and default and node$1 and switch$3 and fabric) and @py<=2001	9
☐	L12	L1 and snmp	0
☐	L11	L10 and L1	1
☐	L10	(pre$post) near5 (buffer$1)	17
☐	L9	L1 and (switched near5 fabric)	4
☐	L8	L1 and datagram	4
☐	L7	L1 and subnet	3
☐	L6	L5 and traffic	3
☐	L5	L4 and node$1	6
☐	L4	L3 and host	11
☐	L3	L1 and messages	32
☐	L2	L1 and (queue near5 pair$1)	10
	L1	(buffer$1 and queue$).ti.	439

☐

END OF SEARCH HISTORY

Terms used
infiniband buffers fabric messages threshold queue pairs management

Found **25,988** of **176,279**

Sort results by relevance

Display results expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown

Relevance scale

1 Communication systems: Software-directed power-aware interconnection networks

Vassos Soteriou, Noel Eisley, Li-Shiuan Peh

September 2005 **Proceedings of the 2005 international conference on Compilers, architectures and synthesis for embedded systems CASES '05**

Publisher: ACM Press

Full text available: pdf(895.11 KB) Additional Information: full citation, abstract, references, index terms

Interconnection networks have been deployed as the communication fabric in a wide range of parallel computer systems. With recent technological trends allowing growing quantities of chip resources and faster clock rates, there have been prevailing concerns of increasing power consumption being a major limiting factor in the design of parallel computer systems, from multiprocessor SoCs to multi-chip embedded systems and parallel servers. To tackle this, power-aware networks must become inherent c ...

Keywords: communication links, dynamic voltage, interconnection networks, networks on-a-chip (NoC), scaling, simulation, software-directed power reduction

2 Promises and reality: Server I/O networks past, present, and future

Renato John Recio

August 2003 **Proceedings of the ACM SIGCOMM workshop on Network-I/O convergence: experience, lessons, implications**

Publisher: ACM Press

Full text available: pdf(225.62 KB) Additional Information: full citation, abstract, references, index terms

Enterprise and technical customers place a diverse set of requirements on server I/O networks. In the past, no single network type has been able to satisfy all of these requirements. As a result several fabric types evolved and several interconnects emerged to satisfy a subset of the requirements. Recently several technologies have emerged that enable a single interconnect to be used as more than one fabric type. This paper will describe the requirements customers place on server I/O networks; t ...

Keywords: 10 GigE, Cluster, Cluster Networks, Gigabit Ethernet, I/O Expansion Network, IOEN, InfiniBand, LAN, PCI, PCI Express, RDMA, RNIC, SAN, Socket Extensions, TOE, iONIC, iSCSI, iSER

3

Viable opto-electronic HPC interconnect fabrics

Ronald Luijten, Cyriel Minkenberg, Roe Hemenway, Michael Sauer, Richard Grzybowski

November 2005 **Proceedings of the 2005 ACM/IEEE conference on Supercomputing SC '05**

Publisher: IEEE Computer Society

Full text available: pdf(561.38 KB) Publisher Site Additional Information: full citation, abstract

We address the problem of how to exploit optics for ultrascale High Performance Computing interconnect fabrics. We show that for high port counts these fabrics require multistage topologies regardless of whether electronic or optical switch components are used. Also, per stage electronic buffers remain indispensable for maintaining throughput, lossless-ness and packet sequence. Although the notion of true all-optical packet switching is not yet viable, we show that appropriate use of optical swi ...

Keywords: HPC, Interconnect, Switching, Optical Switching

4 Queue pair IP: a hybrid architecture for system area networks

Philip Buonadonna, David Culler

May 2002 **ACM SIGARCH Computer Architecture News , Proceedings of the 29th annual international symposium on Computer architecture ISCA '02 , Proceedings of the 29th annual international symposium on Computer architecture ISCA '02**, Volume 30 Issue 2

Publisher: IEEE Computer Society, ACM Press

Full text available: pdf(1.01 MB) Publisher Site Additional Information: full citation, abstract, references, citings, index terms

We propose a SAN architecture called Queue Pair IP (QPIP) that combines the interface from industry proposals for low overhead, high bandwidth networks, e.g. Infiniband, with the well established inter-network protocol suite. We evaluate how effectively the queue pair abstraction enables inter-network protocol offload. We develop a prototype QPIP system that implements basic queue pair operations over a subset of TCP, UDP and IPv6 protocols using a programmable network adapter. We assess this pr ...

Keywords: Interconnection Networks, Network Interfaces, Internetworking, Distributed Computing

5 Xunet 2: lessons from an early wide-area ATM testbed

Charles R. Kalmanek, Srinivasan Keshav, William T. Marshall, Samuel P. Morgan, Robert C. Restrick

February 1997 **IEEE/ACM Transactions on Networking (TON)**, Volume 5 Issue 1

Publisher: IEEE Press

Full text available: pdf(231.69 KB) Additional Information: full citation, references, index terms

Keywords: asynchronous transfer mode, available bit rate, constant bit rate, variable bit rate

6 Performance and fluid simulations of a novel shared buffer management system

Krishnan Kumaran, Debasis Mitra

January 2001 **ACM Transactions on Modeling and Computer Simulation (TOMACS)**, Volume 11 Issue 1

Publisher: ACM Press

Full text available: pdf(229.12 KB) Additional Information: full citation, abstract, references, citings, index terms

We consider a switching system that has multiple ports that share a common buffer, in which there is a FIFO logical queue for each port. Each port may support a large number of flows or connections, which are approximately homogeneous in their statistical characteristics, with common QoS requirements in cell loss and maximum delay. Heterogeneity may exist across ports. Our first contribution is a buffer management scheme based on Buffer Admission Control, which is integrated with Connection ...

Keywords: buffer management, fluid simulations, virtual partitioning

7 Communication: RDMA read based rendezvous protocol for MPI over InfiniBand: design alternatives and benefits

Sayantan Sur, Hyun-Wook Jin, Lei Chai, Dhabaleswar K. Panda

March 2006 **Proceedings of the eleventh ACM SIGPLAN symposium on Principles and practice of parallel programming PPoPP '06**

Publisher: ACM Press

Full text available: pdf(322.07 KB) Additional Information: full citation, abstract, references, index terms

Message Passing Interface (MPI) is a popular parallel programming model for scientific applications. Most high-performance MPI implementations use Rendezvous Protocol for efficient transfer of large messages. This protocol can be designed using either RDMA Write or RDMA Read. Usually, this protocol is implemented using RDMA Write. The RDMA Write based protocol requires a two-way handshake between the sending and receiving processes. On the other hand, to achieve low latency, MPI implementations ...

Keywords: InfiniBand, MPI, communication overlap

8 STORM: lightning-fast resource management

Eitan Frachtenberg, Fabrizio Petrini, Juan Fernandez, Scott Pakin, Salvador Coll

November 2002 **Proceedings of the 2002 ACM/IEEE conference on Supercomputing**

Publisher: IEEE Computer Society Press

Full text available: pdf(324.44 KB) Additional Information: full citation, abstract, references, citings, index terms

Although workstation clusters are a common platform for high-performance computing (HPC), they remain more difficult to manage than sequential systems or even symmetric multiprocessors. Furthermore, as cluster sizes increase, the quality of the resource-management subsystem---essentially, all of the code that runs on a cluster other than the applications---increasingly impacts application efficiency. In this paper, we present STORM, a resource-management framework designed for scalability and pe ...

9 Cluster communication protocols for parallel-programming systems

Kees Verstoep, Raoul A. F. Bhoedjang, Tim Rühl, Henri E. Bal, Rutger F. H. Hofman

August 2004 **ACM Transactions on Computer Systems (TOCS)**, Volume 22 Issue 3

Publisher: ACM Press

Full text available: pdf(1.29 MB) Additional Information: full citation, abstract, references, index terms

Clusters of workstations are a popular platform for high-performance computing. For many parallel applications, efficient use of a fast interconnection network is essential for good performance. Several modern System Area Networks include programmable network interfaces that can be tailored to perform protocol tasks that otherwise would need to be done by the host processors. Finding the right trade-off between protocol processing at

the host and the network interface is difficult in general. In ...

Keywords: Clusters, parallel-programming systems, system area networks

10 A comparison of architectural support for messaging in the TMC CM-5 and the Cray T3D

Vijay Karamcheti, Andrew A. Chien

May 1995 **ACM SIGARCH Computer Architecture News , Proceedings of the 22nd annual international symposium on Computer architecture ISCA '95**, Volume 23 Issue 2

Publisher: ACM Press

Full text available: pdf(1.21 MB) Additional Information: full citation, abstract, references, citings, index terms

Programming models based on messaging continue to be an important programming model for parallel machines. Messaging costs are strongly influenced by a machine's network interface architecture. We examine the impact of architectural support for messaging in two machines --- the TMC CM-5 and the Cray T3D --- by exploring the design and performance of several messaging implementations. The additional features in the T3D support remote operations: memory access, fetch-and-increment, atomic swaps, a ...

11 GPGPU: general purpose computation on graphics hardware

David Luebke, Mark Harris, Jens Krüger, Tim Purcell, Naga Govindaraju, Ian Buck, Cliff Woolley, Aaron Lefohn

August 2004 **Proceedings of the conference on SIGGRAPH 2004 course notes GRAPH '04**

Publisher: ACM Press

Full text available: pdf(63.03 MB) Additional Information: full citation, abstract

The graphics processor (GPU) on today's commodity video cards has evolved into an extremely powerful and flexible processor. The latest graphics architectures provide tremendous memory bandwidth and computational horsepower, with fully programmable vertex and pixel processing units that support vector operations up to full IEEE floating point precision. High level languages have emerged for graphics hardware, making this computational power accessible. Architecturally, GPUs are highly parallel s ...

12 Building Multirail InfiniBand Clusters: MPI-Level Design and Performance Evaluation

Jiuxing Liu, Abhinav Vishnu, Dhabaleswar K. Panda

November 2004 **Proceedings of the 2004 ACM/IEEE conference on Supercomputing**

Publisher: IEEE Computer Society

Full text available: pdf(220.90 KB) Additional Information: full citation, abstract

In the area of cluster computing, InfiniBand is becoming increasingly popular due to its open standard and high performance. However, even with InfiniBand, network bandwidth can still become the performance bottleneck for some of todayýs most demanding applications. In this paper, we study the problem of how to overcome the bandwidth bottleneck by using multirail networks. We present different ways of setting up multirail networks with InfiniBand and propose a unified MPI design that can support ...

13 Storage protocol designs: A study of iSCSI extensions for RDMA (iSER)

Mallikarjun Chadalapaka, Hemal Shah, Uri Elzur, Patricia Thaler, Michael Ko

August 2003 **Proceedings of the ACM SIGCOMM workshop on Network-I/O convergence: experience, lessons, implications**

Publisher: ACM Press

Full text available: Additional Information:

pdf(281.32 KB) full citation, abstract, references, index terms

The iSCSI protocol is the IETF standard that maps the SCSI family of application protocols onto TCP/IP enabling convergence of storage traffic on to standard TCP/IP fabrics. The ability to efficiently transfer and place the data on TCP/IP networks is crucial for this convergence of the storage traffic. The iWARP protocol suite provides Remote Direct Memory Access (RDMA) semantics over TCP/IP networks and enables efficient memory-to-memory data transfers over an IP fabric. This paper studies the ...

Keywords: DA, DDP, DI, Datamover, MPA, RDMA, RDMAP, SCSI, Verbs, iSCSI, iSER, iWARP

14 TinyDB: an acquisitional query processing system for sensor networks

Samuel R. Madden, Michael J. Franklin, Joseph M. Hellerstein, Wei Hong
March 2005 **ACM Transactions on Database Systems (TODS)**, Volume 30 Issue 1
Publisher: ACM Press
Full text available: pdf(1.67 MB) Additional Information: full citation, abstract, references, index terms

We discuss the design of an acquisitional query processor for data collection in sensor networks. Acquisitional issues are those that pertain to where, when, and how often data is physically acquired (*sampled*) and delivered to query processing operators. By focusing on the locations and costs of acquiring data, we are able to significantly reduce power consumption over traditional passive systems that assume the a priori existence of data. We discuss simple extensions to SQL for controlli ...

Keywords: Query processing, data acquisition, sensor networks

15 A hybrid handover protocol for local area wireless ATM networks

Chai-Keong Toh
December 1996 **Mobile Networks and Applications**, Volume 1 Issue 3
Publisher: Kluwer Academic Publishers
Full text available: pdf(960.44 KB) Additional Information: full citation, abstract, references, citings, index terms

While handovers of voice calls in a wide area mobile environment are well understood, handovers of multi-media traffic in a local area mobile environment is still in its early stage of investigation. Unlike the public wireless networks, handovers for multi-media Wireless LANs (WLANs) have special requirements. In this paper, the problems and challenges faced in a multi-media WLAN environment are outlined and a multi-tier wireless cell clustering architecture is introduced. Design issues for ...

16 Experiences with VI communication for database storage

Yuanyuan Zhou, Angelos Bilas, Suresh Jagannathan, Cezary Dubnicki, James F. Philbin, Kai Li
May 2002 **ACM SIGARCH Computer Architecture News , Proceedings of the 29th annual international symposium on Computer architecture ISCA '02 , Proceedings of the 29th annual international symposium on Computer architecture ISCA '02**, Volume 30 Issue 2
Publisher: IEEE Computer Society, ACM Press
Full text available: pdf(1.29 MB) Publisher Site Additional Information: full citation, abstract, references, citings, index terms

This paper examines how VI-based interconnects can be used to improve I/O path performance between a database server and the storage subsystem. We design and implement a software layer, DSA, that is layered between the application and VI. DSA takes advantage of specific VI features and deals with many of its shortcomings. We

provide and evaluate one kernel-level and two user-level implementations of DSA. These implementations trade transparency and generality for performance at different degrees ...

Keywords: Storage system, cluster-based storage, Database storage, storage area network, User-level Communication, Virtual Interface Architecture, processor overhead

17 Estimation of the cell loss ratio in ATM networks with a fuzzy system and application to measurement-based call admission control

Brahim Bensaou, Shirley T. C. Lam, Hon-Wai Chu, Danny H. K. Tsang

August 1997 **IEEE/ACM Transactions on Networking (TON)**, Volume 5 Issue 4

Publisher: IEEE Press

Full text available: pdf(249.06 KB) Additional Information: full citation, references, citings, index terms

Keywords: ATM, call admission control, cell loss ratio, fuzzy logic

18 Computing curricula 2001

September 2001 **Journal on Educational Resources in Computing (JERIC)**

Publisher: ACM Press

Full text available: pdf(613.63 KB) html(2.78 KB) Additional Information: full citation, references, citings, index terms

19 Network Protocols

Andrew S. Tanenbaum

December 1981 **ACM Computing Surveys (CSUR)**, Volume 13 Issue 4

Publisher: ACM Press

Full text available: pdf(3.37 MB) Additional Information: full citation, references, citings, index terms

20 A new buffer management scheme for hierarchical shared memory switches

Abhijit K. Choudhury, Ellen L. Hahne

October 1997 **IEEE/ACM Transactions on Networking (TON)**, Volume 5 Issue 5

Publisher: IEEE Press

Full text available: pdf(197.29 KB) Additional Information: full citation, references, index terms

Keywords: asynchronous transfer mode, backpressure, buffer memories, hierarchical switch, losses, memory management, priorities, pushout, queuing analysis, shared memory systems

Terms used **infiniband buffers** Found **13,980** of **176,279**

Sort results by: relevance
Display results: expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next
Best 200 shown Relevance scale

1 Performance Comparison of MPI Implementations over InfiniBand, Myrinet and Quadrics
Jiuxing Liu, Balasubramanian Chandrasekaran, Jiesheng Wu, Weihang Jiang, Sushmitha Kini, Weikuan Yu, Darius Buntinas, Peter Wyckoff, D K. Panda
November 2003 **Proceedings of the 2003 ACM/IEEE conference on Supercomputing**
Publisher: IEEE Computer Society
Full text available: pdf(407.62 KB) Additional Information: full citation, abstract

In this paper, we present a comprehensive performance comparison of MPI implementations over Infini-Band, Myrinet and Quadrics. Our performance evaluation consists of two major parts. The first part consists of a set of MPI level micro-benchmarks that characterize different aspects of MPI implementations. The second part of the performance evaluation consists of application level benchmarks. We have used the NAS Parallel Benchmarks and the sweep3D benchmark. We not only present the overall perfo ...

2 Simulation and architecture evaluation: Orion: a power-performance simulator for interconnection networks
Hang-Sheng Wang, Xinping Zhu, Li-Shiuan Peh, Sharad Malik
November 2002 **Proceedings of the 35th annual ACM/IEEE international symposium on Microarchitecture**
Publisher: IEEE Computer Society Press
Full text available: pdf(1.14 MB) Publisher Site Additional Information: full citation, abstract, references, citings, index terms

With the prevalence of server blades and systems-on-a-chip (SoCs), interconnection networks are becoming an important part of the microprocessor landscape. However, there is limited tool support available for their design. While performance simulators have been built that enable performance estimation while varying network parameters, these cover only one metric of interest in modern designs. System power consumption is increasingly becoming equally, if not more important than performance. It is ...

3 Communication: RDMA read based rendezvous protocol for MPI over InfiniBand: design alternatives and benefits
Sayantan Sur, Hyun-Wook Jin, Lei Chai, Dhabaleswar K. Panda
March 2006 **Proceedings of the eleventh ACM SIGPLAN symposium on Principles and practice of parallel programming PPoPP '06**

Publisher: ACM Press
Full text available: pdf(322.07 KB) Additional Information: full citation, abstract, references, index terms

Message Passing Interface (MPI) is a popular parallel programming model for scientific applications. Most high-performance MPI implementations use Rendezvous Protocol for efficient transfer of large messages. This protocol can be designed using either RDMA Write or RDMA Read. Usually, this protocol is implemented using RDMA Write. The RDMA Write based protocol requires a two-way handshake between the sending and receiving processes. On the other hand, to achieve low latency, MPI implementations ...

Keywords: InfiniBand, MPI, communication overlap

4 Empirical evaluation of multi-level buffer cache collaboration for storage systems

Zhifeng Chen, Yan Zhang, Yuanyuan Zhou, Heidi Scott, Berni Schiefer
June 2005 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2005 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '05**, Volume 33 Issue 1
Publisher: ACM Press
Full text available: pdf(379.25 KB) Additional Information: full citation, abstract, references, index terms

To bridge the increasing processor-disk performance gap, buffer caches are used in both storage clients (e.g. database systems) and storage servers to reduce the number of slow disk accesses. These buffer caches need to be managed effectively to deliver the performance commensurate to the aggregate buffer cache size. To address this problem, two paradigms have been proposed recently to *collaboratively* manage these buffer caches together: the ***hierarchy-aware caching*** maintains ...

Keywords: collaborative caching, database, file system, storage system

5 Storage protocol designs: A study of iSCSI extensions for RDMA (iSER)

Mallikarjun Chadalapaka, Hemal Shah, Uri Elzur, Patricia Thaler, Michael Ko
August 2003 **Proceedings of the ACM SIGCOMM workshop on Network-I/O convergence: experience, lessons, implications**
Publisher: ACM Press
Full text available: pdf(281.32 KB) Additional Information: full citation, abstract, references, index terms

The iSCSI protocol is the IETF standard that maps the SCSI family of application protocols onto TCP/IP enabling convergence of storage traffic on to standard TCP/IP fabrics. The ability to efficiently transfer and place the data on TCP/IP networks is crucial for this convergence of the storage traffic. The iWARP protocol suite provides Remote Direct Memory Access (RDMA) semantics over TCP/IP networks and enables efficient memory-to-memory data transfers over an IP fabric. This paper studies the ...

Keywords: DA, DDP, DI, Datamover, MPA, RDMA, RDMAP, SCSI, Verbs, iSCSI, iSER, iWARP

6 Internet nuggets: Internet nuggets

Mark Thorson
March 2003 **ACM SIGARCH Computer Architecture News**, Volume 31 Issue 1
Publisher: ACM Press
Full text available: pdf(415.90 KB) Additional Information: full citation, index terms

7 Apex-Map: A Global Data Access Benchmark to Analyze HPC Systems and Parallel Programming Paradigms
Erich Strohmaier, Hongzhang Shan
November 2005 **Proceedings of the 2005 ACM/IEEE conference on Supercomputing SC '05**
Publisher: IEEE Computer Society
Full text available: pdf(406.59 KB) Publisher Site
Additional Information: full citation, abstract

The memory wall and global data movement have become thedominant performance bottleneck for many scientific applications.New characterizations of data access streams and related benchmarks to measure their performances are therefore needed to compare HPC systems,software, and programming paradigms effectively. In this paper, weintroduce a novel global data access benchmark, Apex-Map. It is a parameterized synthetic performance probe and integrates concepts fortemporal and spatial locality into i ...

8 Ultra-high performance communication with MPI and the Sun fire™ link interconnect
Steven J. Sistare, Christopher J. Jackson
November 2002 **Proceedings of the 2002 ACM/IEEE conference on Supercomputing**
Publisher: IEEE Computer Society Press
Full text available: pdf(115.46 KB)
Additional Information: full citation, abstract, references, citings, index terms

We present a new low-latency system area network that provides the ultra-high bandwidth needed to fuse a collection of large SMP servers into a capability cluster. The network adapter exports a remote shared memory (RSM) model that supports low latency kernel bypass messaging. The Sun™ MPI library uses the RSM interface to implement a highly efficient memory-to-memory messaging protocol in which the library directly manages buffers and data structures in remote memory. This allows flexible ...

Keywords: MPI, SAN, interconnects, kernel bypass, performance evaluation, remote shared memory

9 An Application-Based Performance Characterization of the Columbia Supercluster
Rupak Biswas, M. Jahed Djomehri, Robert Hood, Haoqiang Jin, Cetin Kiris, Subhash Saini
November 2005 **Proceedings of the 2005 ACM/IEEE conference on Supercomputing SC '05**
Publisher: IEEE Computer Society
Full text available: pdf(776.50 KB) Publisher Site
Additional Information: full citation, abstract

Columbia is a 10,240-processor supercluster consisting of 20 Altix nodes with 512 processors each, and currently ranked as one of the fastest computers in the world. In this paper, we present the performance characteristics of Columbia obtained on up to four computing nodes interconnected via the InfiniBand and/or NUMAlink4 communication fabrics. We evaluate floatingpoint performance, memory bandwidth, message passing communication speeds, and compilers using a subset of the HPC Challenge benchm ...

Keywords: SGI Altix, multi-level parallelism, HPC Challenge benchmarks, NAS Parallel Benchmarks, molecular dynamics, multi-block overset grids, computational fluid dynamics

10 Poster session 1: 2.5GHz PLL with current matching charge-pump for 10Gbps

transmitter design
Jaehong Ko, Wookwan Lee, Soo-Won Kim
April 2005 **Proceedings of the 15th ACM Great Lakes symposium on VLSI**
Publisher: ACM Press
Full text available: pdf(361.98 KB) Additional Information: full citation, abstract, references, index terms

In this paper, we describe a mixed PLL architecture for low jitter clock generation that use a proposed charge pump and output buffer. The newly designed charge-pump circuit is composed of two differential inputs, the signals of UP and DN from the PFD, and an op-amp that matches between the upper current and the lower current. If the two currents are matched in charge-pump there is a low ripple voltage, which is VCO control voltage. The new charge pump circuit has a good immunity from MOSFET wid ...

Keywords: PLL, charge-pump, jitter, output buffer

11 Storage protocol designs: NFS over RDMA
Brent Callaghan, Theresa Lingutla-Raj, Alex Chiu, Peter Staubach, Omer Asad
August 2003 **Proceedings of the ACM SIGCOMM workshop on Network-I/O convergence: experience, lessons, implications**
Publisher: ACM Press
Full text available: pdf(126.79 KB) Additional Information: full citation, abstract, references

The NFS filesystem was designed as a work-group filesystem, making a central file store available to and shared between a number of client workstations. However, more recently NFS has grown in popularity in the server room, connecting large application servers with back-end file servers. In this environment, where high-speed access to data is critical, high capacity interconnects like gigabit Ethernet, Fibre Channel and Infiniband are to be expected. With RDMA technology we can fully utilize the ...

12 SIGCOMM 2003 conference workshop reports: Workshop on network-I/O convergence: experience, lessons, implications (NICELI)
Vinay Aggarwal, Olaf Maennel, Jeffrey Mogul, Allyn Romanow
October 2003 **ACM SIGCOMM Computer Communication Review**, Volume 33 Issue 5
Publisher: ACM Press
Full text available: pdf(103.04 KB) Additional Information: full citation, abstract

This is a summary of the NICELI workshop, based on scribe reports written by Olaf Maennel and Vinay Aggarwal, and edited by Jeffrey Mogul and Allyn Romanow with help from the NICELI attendees. The workshop was held in conjunction with SIGCOMM 2003 on 27 August 2003 in Karlsruhe, Germany.Papers and presentations from the workshop are available on the Web at http://www.acm.org/sigs/sigcomm/sigcomm2003/workshop/niceli/Note-taking during an active discussion is a fallible process, so these notes may ...

13 STORM: lightning-fast resource management
Eitan Frachtenberg, Fabrizio Petrini, Juan Fernandez, Scott Pakin, Salvador Coll
November 2002 **Proceedings of the 2002 ACM/IEEE conference on Supercomputing**
Publisher: IEEE Computer Society Press
Full text available: pdf(324.44 KB) Additional Information: full citation, abstract, references, citings, index terms

Although workstation clusters are a common platform for high-performance computing (HPC), they remain more difficult to manage than sequential systems or even symmetric multiprocessors. Furthermore, as cluster sizes increase, the quality of the resource-management subsystem---essentially, all of the code that runs on a cluster other than the applications---increasingly impacts application efficiency. In this paper, we present

STORM, a resource-management framework designed for scalability and pe ...

14 Queue pair IP: a hybrid architecture for system area networks

Philip Buonadonna, David Culler
May 2002 **ACM SIGARCH Computer Architecture News , Proceedings of the 29th annual international symposium on Computer architecture ISCA '02 , Proceedings of the 29th annual international symposium on Computer architecture ISCA '02**, Volume 30 Issue 2
Publisher: IEEE Computer Society, ACM Press
Full text available: pdf(1.01 MB) Publisher Site Additional Information: full citation, abstract, references, citings, index terms

We propose a SAN architecture called Queue Pair IP (QPIP) that combines the interface from industry proposals for low overhead, high bandwidth networks, e.g. Infiniband, with the well established inter-network protocol suite. We evaluate how effectively the queue pair abstraction enables inter-network protocol offload. We develop a prototype QPIP system that implements basic queue pair operations over a subset of TCP, UDP and IPv6 protocols using a programmable network adapter. We assess this pr ...

Keywords: Interconnection Networks, Network Interfaces, Internetworking, Distributed Computing

15 Networks III: A near optimal scheduler for switch-memory-switch routers

Adnan Aziz, Amit Prakash, Vijaya Ramachandra
June 2003 **Proceedings of the fifteenth annual ACM symposium on Parallel algorithms and architectures**
Publisher: ACM Press
Full text available: pdf(1.10 MB) Additional Information: full citation, abstract, references, citings, index terms

We present a simple and near optimal randomized parallel scheduling algorithm for scheduling packets in routers based on the *Switch-Memory-Switch* (*SMS*)architecture, which emulates 'output queuing' by using a collection of small memories within the switch to buffer packets, and which forms the basis of the fastest routers in use today. For a router with N inputs and N outputs, our algorithm computes the schedule in $O(\log^* N)$ rounds, where a round is a com ...

Keywords: matching, parallelism, randomization, routers, schedulers

16 Experiences with VI communication for database storage

Yuanyuan Zhou, Angelos Bilas, Suresh Jagannathan, Cezary Dubnicki, James F. Philbin, Kai Li
May 2002 **ACM SIGARCH Computer Architecture News , Proceedings of the 29th annual international symposium on Computer architecture ISCA '02 , Proceedings of the 29th annual international symposium on Computer architecture ISCA '02**, Volume 30 Issue 2
Publisher: IEEE Computer Society, ACM Press
Full text available: pdf(1.29 MB) Publisher Site Additional Information: full citation, abstract, references, citings, index terms

This paper examines how VI-based interconnects can be used to improve I/O path performance between a database server and the storage subsystem. We design and implement a software layer, DSA, that is layered between the application and VI. DSA takes advantage of specific VI features and deals with many of its shortcomings. We provide and evaluate one kernel-level and two user-level implementations of DSA. These implementations trade transparency and generality for performance at different

degrees ...

Keywords: Storage system, cluster-based storage, Database storage, storage area network, User-level Communication, Virtual Interface Architecture, processor overhead

17 Promises and reality: Server I/O networks past, present, and future

Renato John Recio
August 2003 **Proceedings of the ACM SIGCOMM workshop on Network-I/O convergence: experience, lessons, implications**
Publisher: ACM Press
Full text available: pdf(225.62 KB) Additional Information: full citation, abstract, references, index terms

Enterprise and technical customers place a diverse set of requirements on server I/O networks. In the past, no single network type has been able to satisfy all of these requirements. As a result several fabric types evolved and several interconnects emerged to satisfy a subset of the requirements. Recently several technologies have emerged that enable a single interconnect to be used as more than one fabric type. This paper will describe the requirements customers place on server I/O networks; t ...

Keywords: 10 GigE, Cluster, Cluster Networks, Gigabit Ethernet, I/O Expansion Network, IOEN, InfiniBand, LAN, PCI, PCI Express, RDMA, RNIC, SAN, Socket Extensions, TOE, iONIC, iSCSI, iSER

18 InfiniBand and Linux

Roland Dreier
May 2005 **Linux Journal**, Volume 2005 Issue 133
Publisher: Specialized Systems Consultants, Inc.
Full text available: html(16.87 KB) Additional Information: full citation, abstract, index terms

If 120Gb/s isen't fast enough for you, try receiving data without the CPU doing a thing.

19 Micro-architecture techniques in the intel® E8870 scalable memory controller

Fayé Briggs, Suresh Chittor, Kai Cheng
June 2004 **Proceedings of the 3rd workshop on Memory performance issues: in conjunction with the 31st international symposium on computer architecture WMPI '04**
Publisher: ACM Press
Full text available: pdf(435.63 KB) Additional Information: full citation, abstract, references, index terms

This paper describes several selected micro-architectural tradeoffs and optimizations for the scalable memory controller of the Intel E8870 chipset architecture. The Intel E8870 chipset architecture supports scalable coherent multiprocessor systems using 2 to 16 processors, and a point-to-point Scalability Port (SP) Protocol. The scalable memory controller micro-architecture applies a number of micro-architecture techniques to reduce the local & remote idle and loaded latencies. The performance ...

Keywords: distributed coherency, memory latency, scalability, transaction flows

20 Viable opto-electronic HPC interconnect fabrics

Ronald Luijten, Cyriel Minkenberg, Roe Hemenway, Michael Sauer, Richard Grzybowski
November 2005 **Proceedings of the 2005 ACM/IEEE conference on Supercomputing SC '05**
Publisher: IEEE Computer Society

Full text available: pdf(561.38 KB) Additional Information: full citation, abstract

Publisher Site

We address the problem of how to exploit optics for ultrascale High Performance Computing interconnect fabrics. We show that for high port counts these fabrics require multistage topologies regardless of whether electronic or optical switch components are used. Also, per stage electronic buffers remain indispensable for maintaining throughput, lossless-ness and packet sequence. Although the notion of true all-optical packet switching is not yet viable, we show that appropriate use of optical swi ...

Keywords: HPC, Interconnect, Switching, Optical Switching

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Terms used **posting messages buffers** Found **26,872** of **176,279**

Sort results by: relevance
Display results: expanded form
Save results to a Binder
Search Tips
Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next
Best 200 shown Relevance scale

1 A message passing coprocessor for distributed memory multicomputers
Jiun-Ming Hsu, Prithviraj Banerjee
November 1990 **Proceedings of the 1990 ACM/IEEE conference on Supercomputing**
Publisher: IEEE Computer Society
Full text available: pdf(1.25 MB) Additional Information: full citation, abstract, references, citings

This paper presents the architecture, methodology and performance evaluation of a *message passing coprocessor* (MPC) which can accelerate message communication in a distributed memory multicomputer. The MPC is a microprogrammable processor which off-loads the CPU of the burden of communication and speeds up the software processing by directly executing message passing instructions in microcode. It supports process scheduling, message buffer management, and fast buffer copying. The most uni ...

2 Using message passing for distributed programming: proof rules and disciplines
Richard D. Schlichting, Fred B. Schneider
July 1984 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 6 Issue 3
Publisher: ACM Press
Full text available: pdf(1.66 MB) Additional Information: full citation, references, citings, index terms

3 Decoupling synchronization and data transfer in message passing systems of parallel computers
T. Stricker, J. Stichnoth, D. O'Hallaron, S. Hinrichs, T. Gross
July 1995 **Proceedings of the 9th international conference on Supercomputing**
Publisher: ACM Press
Full text available: pdf(1.12 MB) Additional Information: full citation, references, citings, index terms

4 The design and implementation of zero copy MPI using commodity hardware with a high performance network
Francis O'Carroll, Hiroshi Tezuka, Atsushi Hori, Yutaka Ishikawa
July 1998 **Proceedings of the 12th international conference on Supercomputing**
Publisher: ACM Press
Full text available: pdf(803.80 KB) Additional Information: full citation, references, citings, index terms

5 Repeatable and portable message-passing programs

Robert Cypher, Eric Leu
August 1994 **Proceedings of the thirteenth annual ACM symposium on Principles of distributed computing**
Publisher: ACM Press
Full text available: pdf(991.15 KB) Additional Information: full citation, references, citings, index terms

6 A Stochastic Model for Message Assembly Buffering with a Comparison of Block Assignment Strategies

Gary D. Schultz
July 1972 **Journal of the ACM (JACM)**, Volume 19 Issue 3
Publisher: ACM Press
Full text available: pdf(737.75 KB) Additional Information: full citation, references, citings, index terms

7 Storage protocol designs: NFS over RDMA

Brent Callaghan, Theresa Lingutla-Raj, Alex Chiu, Peter Staubach, Omer Asad
August 2003 **Proceedings of the ACM SIGCOMM workshop on Network-I/O convergence: experience, lessons, implications**
Publisher: ACM Press
Full text available: pdf(126.79 KB) Additional Information: full citation, abstract, references

The NFS filesystem was designed as a work-group filesystem, making a central file store available to and shared between a number of client workstations. However, more recently NFS has grown in popularity in the server room, connecting large application servers with back-end file servers. In this environment, where high-speed access to data is critical, high capacity interconnects like gigabit Ethernet, Fibre Channel and Infiniband are to be expected. With RDMA technology we can fully utilize the ...

8 Evaluating the locality benefits of active messages

Ellen Spertus, William J. Dally
August 1995 **ACM SIGPLAN Notices , Proceedings of the fifth ACM SIGPLAN symposium on Principles and practice of parallel programming PPOPP '95**, Volume 30 Issue 8
Publisher: ACM Press
Full text available: pdf(978.43 KB) Additional Information: full citation, abstract, references, citings, index terms

A major challenge in fine-grained computing is achieving locality without excessive scheduling overhead. We built two J-Machine implementations of a fine-grained programming model, the Berkeley Threaded Abstract Machine. One implementation takes an Active Messages approach, maintaining a scheduling hierarchy in software in order to improve data cache performance. Another approach relies on the J-Machine's message queues and fast task switch, lowering the control costs at the expense of data ...

9 Remote queues: exposing message queues for optimization and atomicity

Eric A. Brewer, Frederic T. Chong, Lok T. Liu, Shamik D. Sharma, John D. Kubiatowicz
July 1995 **Proceedings of the seventh annual ACM symposium on Parallel algorithms and architectures**
Publisher: ACM Press
Full text available: pdf(1.78 MB) Additional Information: full citation, references, citings, index terms

10 Realizing the performance potential of the virtual interface architecture

Evan Speight, Hazim Abdel-Shafi, John K. Bennett
May 1999 **Proceedings of the 13th international conference on Supercomputing**
Publisher: ACM Press
Full text available: pdf(1.52 MB) Additional Information: full citation, references, citings, index terms

11 Synchronization for a multi-port frame buffer on a mesh-connected multicomputer

Bin Wei, Gordon Stoll, Douglas W. Clark, Edward W. Felten, Kai Li, Patrick Hanrahan
December 1995 **Proceedings of the IEEE symposium on Parallel rendering**
Publisher: ACM Press
Full text available: pdf(856.27 KB) Additional Information: full citation, citings, index terms

Keywords: mesh routing network, multi-port frame buffer, multicomputers, parallel rendering, synchronization

12 Storage protocol designs: A study of iSCSI extensions for RDMA (iSER)

Mallikarjun Chadalapaka, Hemal Shah, Uri Elzur, Patricia Thaler, Michael Ko
August 2003 **Proceedings of the ACM SIGCOMM workshop on Network-I/O convergence: experience, lessons, implications**
Publisher: ACM Press
Full text available: pdf(281.32 KB) Additional Information: full citation, abstract, references, index terms

The iSCSI protocol is the IETF standard that maps the SCSI family of application protocols onto TCP/IP enabling convergence of storage traffic on to standard TCP/IP fabrics. The ability to efficiently transfer and place the data on TCP/IP networks is crucial for this convergence of the storage traffic. The iWARP protocol suite provides Remote Direct Memory Access (RDMA) semantics over TCP/IP networks and enables efficient memory-to-memory data transfers over an IP fabric. This paper studies the ...

Keywords: DA, DDP, DI, Datamover, MPA, RDMA, RDMAP, SCSI, Verbs, iSCSI, iSER, iWARP

13 Session 13: scheduling and operating systems: Scheduling of unstructured communication on the Intel iPSC/860

Jhy Chun Wang, Sanjay Ranka
November 1994 **Proceedings of the 1994 ACM/IEEE conference on Supercomputing**
Publisher: ACM Press
Full text available: pdf(805.10 KB) Additional Information: full citation, abstract, references

In this paper we present several algorithms for decomposing all-to-many personalized communication into a set of disjoint partial permutations. These partial permutations avoid node contention as well as link contention. We discuss the theoretical complexity of these algorithms and study their effectiveness both from the view of static scheduling and from runtime scheduling. Experimental results for our algorithms are presented on the iPSC/860.

14 Bounds on the efficiency of message-passing protocols for parallel computers

Robert Cypher, Smaragda Konstantinidou

August 1993 **Proceedings of the fifth annual ACM symposium on Parallel algorithms and architectures**
Publisher: ACM Press
Full text available: pdf(1.05 MB) Additional Information: full citation, references, index terms, review

15 An implementation and analysis of the virtual interface architecture
Philip Buonadonna, Andrew Geweke, David Culler
November 1998 **Proceedings of the 1998 ACM/IEEE conference on Supercomputing (CDROM)**
Publisher: IEEE Computer Society
Full text available: html(60.53 KB) Additional Information: full citation, abstract, references, citings

Rapid developments in networking technology and a rise in clustered computing have driven research studies in high performance communication architectures. In an effort to standardize the work in this area, industry leaders have developed the Virtual Interface Architecture (VIA) specification. This architecture seeks to provide an operating system-independent infrastructure for high-performance user-level networking in a generic environment. This paper evaluates the inherent costs and performanc ...

Keywords: cluster, interconnect, network, system-area, user-level, virtual interface architecture

16 The NX/2 operating system
P. Pierce
January 1988 **Proceedings of the third conference on Hypercube concurrent computers and applications: Architecture, software, computer systems, and general issues - Volume 1**
Publisher: ACM Press
Full text available: pdf(598.36 KB) Additional Information: full citation, abstract, citings, index terms

NX/2 is the operating system which runs on the nodes of the Intel iPSC®/2 concurrent supercomputer. NX/2 provides all of the standard system services found in the original iPSC node operating system, such as memory management, multiple process control, message passing services, and intertask protection. This paper focuses on the major node operating system enhancements brought about by two different requirements. First, NX/2 had to support very high speed and high throughput ...

17 Optimal multiphase complete exchange on circuit-switched hypercube architectures
David M. Nicol, Shahid H. Bokhari
May 1994 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1994 ACM SIGMETRICS conference on Measurement and modeling of computer systems SIGMETRICS '94**, Volume 22 Issue 1
Publisher: ACM Press
Full text available: pdf(869.05 KB) Additional Information: full citation, abstract, references, citings, index terms

The complete-exchange communication primitive on a distributed memory multiprocessor calls for every processor to send a message to every other processor, each such message being unique. For circuit-switched hypercube networks there are two well-known schemes for implementing this primitive. Direct exchange minimizes communication volume but maximizes startup costs, while Standard Exchange minimizes startup costs at the price of higher communication volume. ...

18 Communication: RDMA read based rendezvous protocol for MPI over InfiniBand:

design alternatives and benefits
Sayantan Sur, Hyun-Wook Jin, Lei Chai, Dhabaleswar K. Panda
March 2006 **Proceedings of the eleventh ACM SIGPLAN symposium on Principles and practice of parallel programming PPoPP '06**
Publisher: ACM Press
Full text available: pdf(322.07 KB) Additional Information: full citation, abstract, references, index terms

Message Passing Interface (MPI) is a popular parallel programming model for scientific applications. Most high-performance MPI implementations use Rendezvous Protocol for efficient transfer of large messages. This protocol can be designed using either RDMA Write or RDMA Read. Usually, this protocol is implemented using RDMA Write. The RDMA Write based protocol requires a two-way handshake between the sending and receiving processes. On the other hand, to achieve low latency, MPI implementations ...

Keywords: InfiniBand, MPI, communication overlap

19 Compile/run-time support for threaded MPI execution on multiprogrammed shared memory machines
Hong Tang, Kai Shen, Tao Yang
May 1999 **ACM SIGPLAN Notices , Proceedings of the seventh ACM SIGPLAN symposium on Principles and practice of parallel programming PPoPP '99,** Volume 34 Issue 8
Publisher: ACM Press
Full text available: pdf(1.54 MB) Additional Information: full citation, abstract, references, citings, index terms

MPI is a message-passing standard widely used for developing high-performance parallel applications. Because of the restriction in the MPI computation model, conventional implementations on shared memory machines map each MPI node to an OS process, which suffers serious performance degradation in the presence of multiprogramming, especially when a space/time sharing policy is employed in OS job scheduling. In this paper, we study compile-time and run-time support for MPI by using threads and dem ...

20 MPI and Java-MPI: contrasts and comparisons of low-level communication performance
Vladimir Getov, Paul Gray, Vaidy Sunderam
January 1999 **Proceedings of the 1999 ACM/IEEE conference on Supercomputing (CDROM)**
Publisher: ACM Press
Full text available: pdf(96.21 KB) Additional Information: full citation, references, citings, index terms

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

RESULT LIST
6 results found in the Worldwide database for:
infiniband in the title AND **buffers** in the title or abstract
(Results are sorted by date of upload in database)

1 **Infiniband work and completion queue management via head and tail circular buffers with indirect work queue entries**
Inventor: CRADDOCK DAVID F (GB); GREGG THOMAS ANTHONY (US); (+3)
Applicant: IBM (US)
EC: G06F9/46R6M
IPC: ***G06F9/46***; **G06F9/46**; (IPC1-7): G06F13/00 (+1)
Publication info: **TW583544B** - 2004-04-11

2 **Infiniband work and completion queue management via head only circular buffers**
Inventor: ARNDT RICHARD LOUIS (US); CRADDOCK DAVID F (GB); (+3)
Applicant: IBM (US)
EC: H04L12/56Q; H04L12/56Q1
IPC: ***H04L12/56***; **H04L12/56**; (IPC1-7): G06F13/00 (+1)
Publication info: **TW583543B** - 2004-04-11

3 **INFINIBAND CHANNEL ADAPTER FOR PERFORMING DIRECT DMA BETWEEN PCI BUS AND INFINIBAND LINK**
Inventor: PETTEY CHRISTOPHER J; RUBIN LAWRENCE H
Applicant: BANDERACOM INC (US)
EC: G06F13/28; G06F13/38A2; (+1)
IPC: ***G06F13/28; G06F13/38; G06F13/40*** (+6)
Publication info: **WO0235367** - 2002-05-02

4 **Infiniband work and completion queue management via head and tail circular buffers with indirect work queue entries**
Inventor: CRADDOCK DAVID F (US); GREGG THOMAS ANTHONY (US); (+4)
Applicant: IBM (US)
EC: G06F9/46R6M
IPC: ***G06F9/46***; **G06F9/46**; (IPC1-7): G06F3/00
Publication info: **US2003061417** - 2003-03-27

5 **Infiniband work and completion queue management via head only circular buffers**
Inventor: ARNDT RICHARD LOUIS (US); CRADDOCK DAVID F (US); (+5)
Applicant: IBM (US)
EC: H04L12/56Q; H04L12/56Q1
IPC: ***H04L12/56***; **H04L12/56**; (IPC1-7): H04L12/28
Publication info: **US2003058875** - 2003-03-27

6 **Method and apparatus for over-advertising infiniband buffering resources**
Inventor: PEKKALA RICK (US); PETTEY CHRISTOPHER J (US); (+1)
Applicant:
EC: H04L12/56D; H04L12/56E8
IPC: ***H04L12/56***; **H04L12/56**; (IPC1-7): H04L1/00
Publication info: **US2002085493** - 2002-07-04

Data supplied from the ***esp@cenet*** database - Worldwide

RESULT LIST
6 results found in the Worldwide database for:
infiniband in the title AND **buffers** in the title or abstract
(Results are sorted by date of upload in database)

1 **Infiniband work and completion queue management via head and tail circular buffers with indirect work queue entries**
Inventor: CRADDOCK DAVID F (GB); GREGG THOMAS ANTHONY (US); (+3)
Applicant: IBM (US)
EC: G06F9/46R6M
IPC: ***G06F9/46***; **G06F9/46**; (IPC1-7): G06F13/00 (+1)
Publication info: **TW583544B** - 2004-04-11

2 **Infiniband work and completion queue management via head only circular buffers**
Inventor: ARNDT RICHARD LOUIS (US); CRADDOCK DAVID F (GB); (+3)
Applicant: IBM (US)
EC: H04L12/56Q; H04L12/56Q1
IPC: ***H04L12/56***; **H04L12/56**; (IPC1-7): G06F13/00 (+1)
Publication info: **TW583543B** - 2004-04-11

3 **INFINIBAND CHANNEL ADAPTER FOR PERFORMING DIRECT DMA BETWEEN PCI BUS AND INFINIBAND LINK**
Inventor: PETTEY CHRISTOPHER J; RUBIN LAWRENCE H
Applicant: BANDERACOM INC (US)
EC: G06F13/28; G06F13/38A2; (+1)
IPC: ***G06F13/28; G06F13/38; G06F13/40*** (+6)
Publication info: **WO0235367** - 2002-05-02

4 **Infiniband work and completion queue management via head and tail circular buffers with indirect work queue entries**
Inventor: CRADDOCK DAVID F (US); GREGG THOMAS ANTHONY (US); (+4)
Applicant: IBM (US)
EC: G06F9/46R6M
IPC: ***G06F9/46***; **G06F9/46**; (IPC1-7): G06F3/00
Publication info: **US2003061417** - 2003-03-27

5 **Infiniband work and completion queue management via head only circular buffers**
Inventor: ARNDT RICHARD LOUIS (US); CRADDOCK DAVID F (US); (+5)
Applicant: IBM (US)
EC: H04L12/56Q; H04L12/56Q1
IPC: ***H04L12/56***; **H04L12/56**; (IPC1-7): H04L12/28
Publication info: **US2003058875** - 2003-03-27

6 **Method and apparatus for over-advertising infiniband buffering resources**
Inventor: PEKKALA RICK (US); PETTEY CHRISTOPHER J (US); (+1)
Applicant:
EC: H04L12/56D; H04L12/56E8
IPC: ***H04L12/56***; **H04L12/56**; (IPC1-7): H04L1/00
Publication info: **US2002085493** - 2002-07-04

Data supplied from the ***esp@cenet*** database - Worldwide

RESULT LIST
6 results found in the Worldwide database for:
infiniband in the title AND **fabric** in the title or abstract
(Results are sorted by date of upload in database)

1 Method and apparatus for transparent communication between a fibre channel network and an infiniband network
Inventor: LIAO FEEJEN (US); MAKISHIMA DENNIS (US); (+3)
Applicant: BROCADE COMM SYSTEMS INC (US)
EC: H04L29/06
IPC: ***H04L29/06***; *H04L29/08*; **H04L29/06** (+2)
Publication info: **US2004024905** - 2004-02-05

2 METHODOLOGY AND MECHANISM FOR REMOTE KEY VALIDATION FOR NGIO/INFINIBAND APPLICATIONS
Inventor: PARTHASARATHY BALAJI (US); LEITNER BRIAN M (US); (+1)
Applicant: INTEL CORP (US); PARTHASARATHY BALAJI (US); (+2)
EC: H04L29/08A7
IPC: ***H04L29/06***; **H04L29/06**; (IPC1-7): H04L12/00
Publication info: **WO02078254** - 2002-10-03

3 Method and apparatus enabling both legacy and new applications to access an InfiniBand fabric via a socket API
Inventor: YADAV SATYENDRA (US)
Applicant:
EC: H04L29/08A; H04L29/12A
IPC: ***H04L29/08; H04L29/12***; *H04L29/06* (+6)
Publication info: **US2003061402** - 2003-03-27

4 Apparatus amd method for disparate fabric data and transaction buffering within infiniband device
Inventor: PEKKALA RICHARD E (US); PETTEY CHRISTOPHER J (US); (+2)
Applicant:
EC: H04L12/56E3
IPC: ***H04L12/56***; **H04L12/56**; (IPC1-7): H04L12/50
Publication info: **US2002172195** - 2002-11-21

5 Infiniband TM work queue to TCP/IP translation
Inventor: PETTEY CHRISTOPHER J (US)
Applicant: BANDERACOM (US)
EC: H04L29/06
IPC: ***H04L29/06***; **H04L29/06**; (IPC1-7): G06F15/16
Publication info: **US2003014544** - 2003-01-16

6 Method and device for LAN emulation over infiniband fabrics
Inventor: DAVIS ARLIN R (US)
Applicant:
EC: H04L12/46B; H04L29/12A
IPC: ***H04L12/46; H04L29/12***; **H04L12/46** (+3)
Publication info: **US2004213220** - 2004-10-28

Data supplied from the ***esp@cenet*** database - Worldwide

RESULT LIST
23 results found in the Worldwide database for:
messages in the title AND **buffers** in the title or abstract
(Results are sorted by date of upload in database)

1 **REDUCING NUMBER OF WRITE OPERATIONS RELATIVE TO DELIVERY OF OUT-OF-ORDER RDMA SEND MESSAGES**
Inventor: BIRAN GIORA (IL); MACHULSKY ZORIK (IL); (+1) Applicant: IBM (US); BIRAN GIORA (IL); (+2)
EC: H04L29/06 IPC: ***G06F15/167***; **G06F15/16**
Publication info: **WO2005060579** - 2005-07-07

2 **SYSTEM, APPARATUS, AND METHOD FOR AUTOMATED HANDLING OF MESSAGES IN TERMINALS**
Inventor: TORVINEN MARKO Applicant: NOKIA CORP (FI); NOKIA INC (US)
EC: H04W2/06; H04L12/58; (+2) IPC: ***H04L12/58; H04Q7/22***; **H04L12/58** (+2)
Publication info: **WO2004099895** - 2004-11-18

3 **In-order delivery of plurality of RDMA messages**
Inventor: BIRAN GIORA (IL); MACHULSKY ZORIK (IL); (+1) Applicant: IBM (US)
EC: IPC: ***G06F15/16***; **G06F15/16**; (IPC1-7): G06F15/16
Publication info: **US2005144310** - 2005-06-30

4 **Efficient zero copy transfer of messages between nodes in a data processing system**
Inventor: BENDER CARL A (US); CARROLL WALKER B (US); (+3) Applicant: IBM (US)
EC: IPC: ***G06F15/16***; **G06F15/16**; (IPC1-7): G06F15/16
Publication info: **US2005091383** - 2005-04-28

5 **Mechanism for managing incoming data messages in a cluster**
Inventor: PINTO OSCAR P (US); SHAH RAJESH R (US) Applicant:
EC: H04L12/56Q1 IPC: ***G06F7/00***; **G06F7/00**; (IPC1-7): G06F7/00
Publication info: **US2003101158** - 2003-05-29

6 **Method and apparatus for processing chain messages (SGL chaining)**
Inventor: JOHNSON STEPHEN B (US); HOGLUND TIMOTHY E (US); (+1) Applicant: LSI LOGIC CORP (US)
EC: IPC: ***G06F13/14; G06F13/38***; **G06F13/14** (+3)
Publication info: **US6810448** - 2004-10-26

7 **APPARATUS AND METHOD FOR EXTRACTING MESSAGES FROM A DATA STREAM**
Inventor: BROSEY STEVEN A Applicant: GEN INSTRUMENT CORP (US)
EC: H04N7/24T2M IPC: ***H04N7/24***; **H04N7/24**; (IPC1-7): H04N7/24
Publication info: **WO0126380** - 2001-04-12

8 **End-of-message handling and interrupt generation in a CAN module providing hardware assembly of multi-frame CAN messages**
Inventor: BIRNS NEIL (NL); SLIVKOFF WILLIAM J (NL); (+3) Applicant: KONINKL PHILIPS ELECTRONICS NV (NL)
EC: G06F15/78; H04L12/40; (+5) IPC: ***G06F15/78; H04L12/40; H04L12/413*** (+13)
Publication info: **EP1085722** - 2001-03-21

9 **Communication system has buffer for non-time-critical messages allocated lower priority by processing device**
Inventor: HEISING ANDREAS (DE); UECKER RAINER (DE); (+1) Applicant: SIEMENS AG (DE)
EC: H04Q3/64 IPC: ***H04Q3/64***; **H04Q3/64**; (IPC1-7): H04M3/38 (+1)

Publication info: **DE10039426** - 2002-02-28

10 system for controlling message transfer between objects having a controller that assumes standby state and without taking out excess messages from a queue

Inventor: KOMINE HIROAKI (JP); YOKOSHI NORIYUKI (JP); (+1) Applicant: FUJITSU LTD (JP)

EC: IPC: ***G06F9/44; G06F9/46***; **G06F9/44** (+2)

Publication info: **US6848107** - 2005-01-25

Data supplied from the ***esp@cenet*** database - Worldwide